

1 **SCHEDULING SEQUENTIAL DATA**
2 **PREFETCHES IN A PREEXISTING LRU CACHE**

3
4
5 **BACKGROUND OF THE INVENTION**

6 **1. The Field of the Invention**

7 The present invention relates to cache management in data storage systems. More
8
9 specifically, the present invention relates to prefetch scheduling in a preexisting LRU cache
10 of a data storage system.

11
12
13 **2. The Relevant Art**

14 Cache memory is used in data storage systems to buffer frequently accessed data in
15 order to allow the data to be accessed at a relatively high rate. The cache memory is a
16 relatively small high speed memory operating on a host processor or between the host
17 processor and relatively slower memory devices. Typical data storage systems using caching
18 may include a cache directory or index of the data elements in a main memory of the hosts
19 operating on the data storage system. The cache directory is referenced to provide an
20 indication of whether or not each data element of the main memory resides in the cache
21 memory at any give time, and if so, to indicate the present location of the data element in the
22 cache memory. When a host processor requests an Input/Output (I/O) operation, the cache
23 directory is first consulted to determine whether the requested data element is present in the
24 cache memory and if so, to determine its location. When the data element is present in the
25
26

1 cache memory, the data element can be quickly accessed, rather than having to be requested
2 from a slower storage device.

3
4 Generally, in such systems, every time a data element is requested, a determination
5 is made whether the accessed data is likely to be accessed again in the near future. If so, the
6 accessed data element is copied or "staged" into the cache memory. In some data storage
7 systems, requested data elements are always staged into the cache memory if they are absent
8 from the cache memory. Some data storage systems are also responsive to explicit "prefetch"
9 commands from the host computer to cause specified data to be staged into the cache, even
10 though the specified data is not immediately accessed by the host computer.
11

12
13 Because the cache memory has a capacity that is smaller than the main memory, it
14 is frequently necessary for data elements in the cache memory to be replaced or removed from
15 the cache memory in order to provide space in the cache memory for more recently requested
16 data elements. In general, for the cache memory to be useful, the data elements removed or
17 replaced from the cache memory must be calculated to be less likely to be accessed in the near
18 future than the new data elements being staged into the cache memory at the time the removal
19 or replacement occurs.
20

21
22 Data storage systems that use disk drives for the main memory typically use random
23 access memory (RAM) for the cache memory. In such a data storage system, the data
24 elements in the cache memory are often logical tracks of data on the disks, although in many
25 systems, the data records are blocks or records of data. The cache directory includes a
26 directory entry for at least each data element stored in the cache. Each directory entry for
each data element stored in the cache memory generally includes a pointer to the location of

1 the data element in the cache memory. The cache directory can be a table including an entry
2 for each data element stored in the disk storage. Alternatively, the directory may include a
3 hash table for accessing lists of the directory entries so that the cache directory need not
4 include any cache directory entries for data elements that are absent from the cache memory.
5 In either case, any one of a plurality of data elements in the cache memory may be replaced
6 or removed from the cache according to the particular cache management scheme being used
7 to make room for another data element.
8

9
10 The performance of such a data storage system is highly dependent on the cache
11 management scheme used for selecting the data element to be removed or replaced. The
12 cache management scheme is implemented by a cache management system, or "cache
13 manager," in the data storage system.
14

15 In one common cache management scheme, a cache manager is programmed to
16 remove or replace the "least-recently-used" (LRU) data element in the cache memory. The
17 least-recently-used data element is usually the data element accessed least recently by the host
18 computer. The cache manager maintains an ordered list, or queue, of the data elements in
19 the cache memory so that the cache manager can readily identify the least-recently-used data
20 element. The queue is typically maintained in a doubly-linked list. When a data element is
21 accessed, the data element is moved to the head of the queue, unless the data element is
22 already at the head of the queue. This process is known as making the data element "young"
23 in the cache and ensures that, when the queue is not empty, the least-recently-used data
24 element in the cache memory will be located at the end of the queue and the most-recently-
25 used element in the cache memory will be located at the head of the queue.
26

1 Data that is fetched into the cache memory may be described in two broad fashions.
2 The first is random access data which denotes data that is needed for specific operations but
3 which is not connected with other data in any manner. Many caching systems are configured
4 for optimal performance when fetching random access data. The second type of data is
5 known as sequential data, denoting that several elements of the data are used by a processor
6 in a specific sequence, typically the sequence in which the data elements are stored on a
7 storage device. Many systems that employ a dedicated or "native" LRU cache are only
8 designed to store data accessed in random access operations and make no provision for
9 accessing sequential data.
10
11

12 Attempts have been made to improve the performance of a native LRU cache when
13 fetching sequential data. These solutions, however, require a modification in some fashion
14 of the LRU cache itself to achieve satisfactory performance for sequential data prefetches.
15 One example of these types of modifications is the creation of a "microcache" within the
16 existing LRU cache to hold the sequential data.
17
18

19 Modifying existing LRU caches is not always a plausible solution. For instance, in
20 legacy systems it may not be possible or desirable to modify the replacement algorithm of the
21 cache. The cache logic or controller may be inaccessible or hardwired, or the system the
22 cache resides on may have been provided for a specific purpose that modifying the cache
23 would disrupt.
24
25

26 Accordingly, a need exists in the art for a method of scheduling prefetches of
sequential data into a native LRU cache without directly modifying the algorithm or structure
of the LRU cache.

OBJECTS AND BRIEF SUMMARY OF THE INVENTION

The data prefetch scheduling system and method of the present invention has been developed in response to the present state of the art, and in particular, in response to the problems and needs in the art that have not yet been fully solved by currently available cache management systems. Accordingly, it is an overall object of the present invention to provide a data prefetch scheduling system and method that overcomes many or all of the above-discussed shortcomings in the art.

To achieve the foregoing object, and in accordance with the invention as embodied and broadly described herein in the preferred embodiment, an improved data prefetch scheduling system and corresponding method are provided. The prefetch scheduling system of the present invention allows for prefetching a stream of sequential data based upon the expected residency time/occupancy time of objects within an existing native least-recently-used (LRU) cache system.

Under the method of the present invention, a stream of Input/Output (I/O) requests between a host and a cache is intercepted, and requested data elements are examined. If logically successive data elements of the I/O stream are not already resident within the LRU cache, the prefetch scheduling system may selectively prestage data elements into the LRU cache. The expected contents of the preexisting LRU cache is tracked, and the information is used to quantify the expected value of prefetching a given data element.

In determining whether to prestage a data element, the requested data element is assigned a priority value based upon its likelihood to be sequentially accessed, or used by the host while within the cache. The priority value is assigned in one embodiment based upon

1 the number of logically preceding data elements present in the cache according to the model
2 of the cache. The assigned priority value is compared against a threshold value to determine
3 if the requested data element is to be prefetched. If the value assigned to the data requested
4 is greater than the threshold value , the prefetch scheduling system schedules one or more
5 prefetches of logically successive data elements.
6

7 Scheduling a prefetch may comprise sending an Input/Output (I/O) request for the
8 logically successive data element to the preexisting LRU cache, which then loads the
9 successive data element into the cache. Because the successive data element was unsolicited
10 by the host processor, the data element is ignored by the host processor.
11

12 The threshold value is preferably dynamic and is adjusted periodically as needed,
13 according to the history of prefetched data elements within the cache. If prefetched data
14 elements have a history of being hit more than other data elements within the cache, the
15 threshold value is decremented, and if prefetched data elements are falling out of the cache
16 without being hit at a greater rate than other elements, the threshold value is incremented.
17
18
19
20
21
22
23
24
25
26

1. *Phragmites australis* (Cav.) Trin. ex Steud. (Common reed)
 2. *Scirpus americanus* (L.) Pers. (Common sedge)
 3. *Scirpus setaceus* (L.) Pers. (Common sedge)
 4. *Scirpus robustus* (L.) Pers. (Common sedge)
 5. *Scirpus cespitosus* (L.) Pers. (Common sedge)
 6. *Scirpus hololepis* (L.) Pers. (Common sedge)
 7. *Scirpus atrovirens* (L.) Pers. (Common sedge)
 8. *Scirpus erioporus* (L.) Pers. (Common sedge)
 9. *Scirpus stramineus* (L.) Pers. (Common sedge)
 10. *Scirpus tabernaemontani* (L.) Pers. (Common sedge)
 11. *Scirpus maritimus* (L.) Pers. (Common sedge)
 12. *Scirpus maritimus* (L.) Pers. (Common sedge)
 13. *Scirpus maritimus* (L.) Pers. (Common sedge)
 14. *Scirpus maritimus* (L.) Pers. (Common sedge)
 15. *Scirpus maritimus* (L.) Pers. (Common sedge)
 16. *Scirpus maritimus* (L.) Pers. (Common sedge)
 17. *Scirpus maritimus* (L.) Pers. (Common sedge)
 18. *Scirpus maritimus* (L.) Pers. (Common sedge)
 19. *Scirpus maritimus* (L.) Pers. (Common sedge)
 20. *Scirpus maritimus* (L.) Pers. (Common sedge)

BRIAN C. KUNZLER
ATTORNEY AT LAW
10 WEST 100 SOUTH, SUITE 425
SALT LAKE CITY, UTAH 84101

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25
- 26

Figure 5 is a schematic flow chart diagram illustrating one embodiment of a method for determining and updating a threshold value of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 is a schematic block diagram illustrating a computer system 10 in which executable and operational data, operating in accordance with the present invention, may be hosted on one or more computer stations 12 in a network 14. The network 14 preferably comprises a storage area network (SAN) but may also comprise a wide area network (WAN) or local area network (LAN) and may also comprise an interconnected system of networks, one particular example of which is the Internet and the World Wide Web supported on the Internet.

A typical computer station 12 may include a processor or CPU 16. The CPU 16 may be operably connected to one or more memory devices 18. The memory devices 18 are depicted as including a non-volatile storage device 20 such as a hard disk drive or CD-ROM drive, a read-only memory (ROM) 22, and a random access volatile memory (RAM) 24. Preferably, the computer station 12 operates under the control of an operating system (OS) 25, such as MVS®, OS/390®, AIX®, OS/2®, WINDOWS NT®, WINDOWS®, UNIX®, and the like.

The computer station 12 or system 10 in general may also include one or more input devices 26, such as a mouse or keyboard, for receiving inputs from a user or from another device. Similarly, one or more output devices 28, such as a monitor or printer, may be provided within or be accessible from the computer system 10. A network port such as a network interface card 30 may be provided for connecting to outside devices through the network 14. In the case where the network 14 is remote from the computer station, the

1 The stations 12 connected on the network 14 may comprise data storage servers 46
2 and/or data storage devices 45 to which may be connected an existing least-recently-used
3 (LRU) cache 42. In the depicted embodiment, the cache 42 is a stand-alone module remote
4 to both the station 12 and the server 46, but, of course, could be implemented within a station
5 12 or a server 46, including a storage server. Other resources or peripherals 44, such as
6 printers and scanners may also be connected to the network 36. Other networks may be in
7 communication with the network 14 through a router 38 and/or over the Internet 40.
8
9

10 Figure 2 is a schematic block diagram illustrating one embodiment of a prefetch
11 module 200, suitable for use in a data prefetch scheduling system. In accordance with the
12 invention, the data prefetch scheduling system may also include a station 12, a cache 42, and
13 a server 46 of Figure 1. In one embodiment, sequential data stored in a storage device 45 of
14 the server 46 is prestaged into the cache 42 of Figure 1 using the prefetch module 200. The
15 prefetch module 200 may reside anywhere in the computer system 10, and in a preferred
16 embodiment resides as a daemon on the computer station 12 of Figure 1. More preferably,
17 however, the prefetch module 200 is located on or between the computer station 12 (the
18 "host") and the data storage server 46 of Figure 1, and most preferably, the prefetch module
19 200 operates on a processor of the computer station 12. The prefetch module 200, as
20 depicted, is configured with an interface module 202, a calculation module 204, a dynamic
21 threshold optimization module 206, a prefetch request module 208, and a remote modeling
22 module 210.
23
24
25
26

The remote modeling module 210 preferably models the operation, and to the extent
known, the contents of the cache 42 of Figure 1. The data prefetch module 200 preferably

Each of the objects 230 preferably contains information about an individual data element of the cache 42 of Figure 1. In one embodiment, each object 230 stores a header 230 identifying the data element modeled by the object 230. Each object 230 may also store a history 234 of the represented data element. In one embodiment, the history 234 is represented by a priority value 236 assigned to the data element. A marker 238 indicating whether the data element 223 was stored as a result of a prefetch operation or not may also be stored within the object 230. A time stamp 240 indicating when a data element first entered the cache is also preferably present within each object 230. Of course, other data that may be needed to accurately model each data element that resides within the cache 42 of Figure 1 may likewise be stored in or with the objects 230.

- Page 11 -

BRIAN C. KUNZLER
ATTORNEY AT LAW
110 WEST 100 SOUTH, SUITE 425
SALT LAKE CITY, UTAH 84101

[illegible]

Sub
A1

1 The objects 230 modeled as being within the cache according to the model 220 are
2 linked to the remote model 220 with a data structure, such as a pointer. Within the model
3 220 or associated with the model 220 is a calculated single reference residency time (SRRT)
4 225 for the cache. The SRRT 225 is used to determine when an object has fallen out of the
5 cache. For instance, in one embodiment, when the time stamp 240 for an object 230 is found
6 to exceed the SRRT of the cache 42, that object 230 is removed from the model 220 of the
7 cache. A counter 222 may also be present within the model 220 to store hit rates, as will be
8 described below.
9

11 The interface module 202 is shown provided with an intercept module 205. The
12 intercept module 205 is preferably programmed or otherwise configured to receive either
13 periodic statistical data or to receive statistical data periodically from the cache 42 of Figure
14 1. This data may comprise the hit rate of the cache 42 of Figure 1, the cache size of the cache
15 42 of Figure 1, and the Input/Output (I/O) rate of the cache 42 of Figure 1. This data is then
16 passed to a SRRT module 215 of the calculation module 204. The SRRT module 215
17 calculates the SRRT 225 For the model 220. The calculation module 204 is also preferably
18 provided with a comparison module 212 that is programmed or otherwise configured to
19 receive the priority value 236 of a data element and to compare that priority value 236 to a
20 threshold value 214 to determine whether the priority value is greater than the threshold value
21 214. This determination is then used to determine whether to prefetch the data element, as
22 will be explained.
23
24
25
26

The intercept module 205 is preferably configured to monitor the transmission of
data between the host 12 and the cache 42 and to intercept I/O requests. Other I/O requests

1 to the cache 42 from other stations 12 may not be intercepted. The intercepted I/O requests
2 are then examined for data elements requested to be transferred. Briefly, when such a data
3 element is requested, the calculation module confers with the remote modeling module 210
4 to determine how many logically adjacent preceding data elements are present within the
5 cache, and if sufficient preceding data elements are present, notifies the prefetch request
6 module 208 to schedule a prefetch of the next n logically successive data elements. The
7 variable n is the numeral 1 in one embodiment, but could be any number and is in one
8 embodiment experimentally determined for optimal operation of the cache 42.

11 In one embodiment, the prefetch request module 208 requests the prefetch by issuing
12 an artificial I/O command that is not solicited by the host processor 16 and as such is ignored
13 by the processor 16. The cache 42, however, is "fooled" into believing that a real I/O request
14 has been made and fetches the data element from the storage device 45, keeping a copy in the
15 cache 42.

18 In one embodiment, the comparison module 212 compares a priority value 236 of
19 the requested data element and compares it to a threshold value 214 to make the
20 determination whether it is likely that successive data elements will be requested by the host
21 12. The priority value 236 is, in one embodiment, calculated according to the number of
22 logically adjacent preceding data elements that are present in the model 220 of the cache. The
23 threshold value is first calculated and then continually optimized by the dynamic threshold
24 optimization module 206 in a manner that will be described below.

26 In a second, more preferred embodiment, a plurality of models 220 are maintained
within the remote modeling module 210. These models in one embodiment comprise a minus

1 1 model 224, a baseline model 226, and a plus 1 model 228. In the baseline model 226, the
2 operation of the cache 42 is modeled using the current threshold value 214. In the minus 1
3 model 224, the cache is modeled using the current threshold value minus one whole value.
4 In the plus 1 model, the cache is modeled using the current threshold value plus one.
5 Periodically, as will be explained, the dynamic threshold optimization module 206 compares
6 the operation of the three models 224, 226, 228 and updates the threshold value 214
7 according to which of the three values is at the time providing the most optimal performance
8 of the cache.
9

10
11 Figure 3 is a schematic flow chart diagram illustrating one embodiment of a data
12 prefetch scheduling method 300 of the present invention. The method 300 in one
13 embodiment may be considered to be the method of use of the prefetch module 200 of Figure
14 2, but may also be used independent of the prefetch module 200 of Figure 2. In one
15 embodiment, the method 300 operates with three concurrent operations, denoted at 301, 309,
16 and 319. The method 300, may upon initialization, step through each step 302 - 334 in
17 sequence, and after that, the separate operations may loop separately.
18

19
20 The data prefetch scheduling method 300 preferably remotely models the
21 performance of the LRU cache 42 of Figure 1, and in so doing the interface module preferably
22 periodically retrieves performance and statistical data from the LRU cache 42 of Figure 1.
23 In a step 302, the size of the LRU cache 42 is determined and is preferably measured in units
24 of I/O requests. The size of the LRU cache 42 is preferably fetched from the LRU cache 42
25 through the interface module 202, but may be determined in other manners such as
26 preprogramming or direct user entry. A step 301 indicates that after the first iteration of the

1 steps 302 - 308, the operation 309 loops at intervals such as every x minutes, where x may
2 be any number and is in one embodiment one minute. Of course, if the cache size is static,
3 step 302 need be conducted only one time upon startup.
4

5 The Input/Output (I/O) rate of the cache 42 of Figure 1 is fetched from the cache 42
6 of Figure 1 by the interface module 204 of Figure 2 in a step 304. The I/O rate refers to the
7 number of I/O requests sent to the cache 42 of Figure 1 by one or more computer stations 12
8 of Figure 1 in a given period of time. The I/O rate includes both hits and "misses" in the
9 cache. The hit rate of the cache 42 of Figure 1 is also preferably fetched from the cache 42
10 of Figure 1 by the interface module 202 of Figure 2 in a step 306. The hit rate refers to the
11 number of I/O requests, i.e., "hits," that were found to reside in the cache at the time of
12 the I/O request over a given period of time. The hit rate may also be used to determine a miss
13 ratio for the cache where the miss ratio is one minus the hit rate divided by the number of I/O
14 requests.
15
16
17

18 The hit rate, in the preferred embodiment, is calculated based upon a logarithmic rate
19 of increase with increased cache size. Calculating the hit rate from the cache size in this
20 manner is necessary, because hit rates do not typically increase in a linear manner with an
21 increasing cache size. Thus, a logarithmic rate of increase is considered a close
22 approximation of the hit rate v. cache size, where the hit rate is estimated to increase at an
23 increasingly lower rate with the increase in cache size.
24
25

26 Once the cache size, I/O rate, and hit rate are fetched, the calculation module 204 of
Figure 2 preferably uses this data to calculate the SRRT 216, which is an approximation of
the average time it takes a data element that has been fetched into the cache to "fall through"

1 or be removed from the cache 42 of Figure 1. The SRRT 216 is later used at a step 322 to
2 determine which data elements have "fallen out of the cache."

3
4 In a preferred embodiment, the SRRT is calculated by estimating the hit rate for all
5 cache sizes 1 to n , where n is the size of the cache as determined in the step 302, and is
6 adjusted downward by the portion of the cache assumed to be holding prefetched requests.
7 Then, using the estimated hit rate of the different cache sizes, the miss ratio for all cache sizes
8 1 to n is calculated. This is done by assuming that the effectiveness of the cache decays
9 exponentially in relation to the size of the cache. Subsequently, using the I/O rate of the
10 cache, the miss ratio of the cache, and the expected miss ratio for the cache one size smaller,
11 the SRRT is calculated using Formula 1,
12
13

$$\frac{1}{r} \left(1 + p + \frac{n-1}{m(n-1)} \right)$$

14
15
16
17
18 Formula 1

19 where r is the I/O rate of the cache, p is the number of prefetches conducted by the
20 prefetching algorithm, n is the size of the cache as determined in the step 302, and $m[i]$ is the
21 expected miss ratio of the cache when the cache is of size i . Of course, other methods exist
22 for calculating the SRRT for an LRU cache and may be used in place of Formula 1. The steps
23 304, 306 and 308 of the data prefetch scheduling method 300 are preferably performed
24 periodically as indicated by the step 301. In the embodiment described above, where a
25 plurality of model 220 is maintained, the SRRT that would occur under operation of each
26

Figure 5 illustrates in greater detail, one embodiment of a method for conducting the threshold value update operation 309 of Figure 3. The method of figure 5 begins at step 400, and is conducted every γ I/O request, as indicted by step 310 on figure 3. At step 401, the

1 hit rate counter 222 of the baseline model 226 of Figure 2 is retrieved. At a step 402, the
2 baseline model counter is compared with the hit rate counter 222 of the minus 1 model 224.
3
4 At a step 403, the results are weighted. That is, the results may be averaged over every *m*
5 iterations of the operation 309. At a step 404, the results are observed. Thus, after the
6 weighting, if the baseline counter is determined to be less than the minus 1 counter, the
7 dynamic threshold is decremented by one integer value at a step 406. The operation 309 then
8 returns to step 310 as indicated at 407. If the results of the observation of step 404 are
9 negative, the baseline counter is compared with the hit rate counter of the plus one model 228
10 at a step 408. Similarly, the results are weighted at a step 409. If, in the weighted results,
11 the baseline counter is less than the plus 1 counter, the dynamic threshold is incremented as
12 indicated at a step 412, and the operation 309 returns back to the step 310 as indicated at
13 413. If the result of the query of step 10 is no, the operation 309 also returns directly to step
14 310 as indicated at 414.
15
16
17

18 Once a threshold value has been selected and/or optimized, the method 300 proceeds
19 to the third operation 319. The operation 319 is also initially conducted in sequence with the
20 operations 301 and 309 and then is conducted independently as data requests are intercepted.
21
22 At a step 320, the operation 319 waits for the intercept module 205 to intercept an I/O
23 request. When an I/O request is intercepted, a data element requested from the storage
24 device 45 or other source is extracted. The models 220 are then updated at a step 322. This
25 preferably comprises comparing the timestamps of the objects 230 within each of the models
26 220 to determine which have exceeded the SRRT for that model. Those exceeding the SRRT
are removed from the model as being likely to have been removed from the cache 42 under

1 that model. Of course, the step 322 could be conducted at any time, and could be a separate
2 independent operation. In one embodiment, this step is repeated later as indicated by step
3 334.
4

5 At a step 324, the hit rates are updated. In one embodiment, this comprises
6 examining the data element identified at the step 320 and comparing that data element to each
7 of the models 220. If the data element is within a model and is marked with the marker 238
8 as having been prefetched, the counter 222 of that model is incremented.
9

10 At a step 326, the models 224, 226, 228 are examined to see if the data element
11 logically preceding the intercepted data element is present within each model. At a step 328,
12 it is determined whether a prefetch of the successive data element(s) is to be conducted
13 according to each of the models 220 and the models 220 are separately updated according
14 to the results of the determination. One embodiment of a method of achieving this process
15 is illustrated in greater detail in Figure 4.
16
17

18 Referring to Figure 4, at a step 350 the method begins, and as indicated at a step 352,
19 is conducted for each of the models 220, beginning with the minus 1 model 224. At a step
20 354, the data element is received into the prefetch module as an object 230. At a step 355,
21 an object 230 is created or updated, depending on whether the data element is already in the
22 cache. In so doing, a timestamp 240 is applied and a header 232 is assigned. The marker 238
23 is also set to indicate whether the data element was prefetched or not. The data element is
24 modeled as being the youngest element in the cache 42, as indicated at a step 356, in
25 accordance with how the cache 42, being a LRU cache, literally treats the data element.
26

1 the priority valve assigned to the data element greater than the threshold valve? If so, at a
2 step 332, the prefetch is scheduled. Scheduling the prefetch preferably comprises passing an
3 unsolicited I/O request to the cache 42, as discussed above.
4

5 In accordance with the present invention as described above, it should be apparent
6 that sequential prefetching of data in a preexisting and/or dedicated LRU cache can now be
7 conducted without significant likelihood of the prefetched data falling out of the cache before
8 being re-referenced. Additionally, this can be accomplished with minimum interference to
9 the "normal" operation of the cache and without having to internally modify the LRU cache
10 or the native operation of the LRU cache. Accordingly, such caches can be enhanced for
11 greater performance with the use of the method and system of the present invention merely
12 with, in one embodiment, the addition of software to a host or between the cache and the
13 host.
14
15

16 The present invention may be embodied in other specific forms without departing
17 from its spirit or essential characteristics. The described embodiments are to be considered
18 in all respects only as illustrative and not restrictive. The scope of the invention is, therefore,
19 indicated by the appended claims rather than by the foregoing description. All changes which
20 come within the meaning and range of equivalency of the claims are to be embraced within
21 their scope.
22
23

24 What is claimed is:
25
26